arXiv:1609.02226v3 [cs.AI] 4 Jan 2017

Fitted Learning: Models with Awareness of their Limits

Navid Kardan and **Kenneth O. Stanley**
Department of Computer Science, University of Central Florida,
4328 Scorpius Street,
Orlando, FL 32816-2362,
kardan@knights.ucf.edu, kstanley@cs.ucf.edu

Abstract

Though deep learning has pushed the boundaries of classification forward, in recent years hints of the limits of standard classification have begun to emerge. Problems such as fooling, adding new classes over time, and the need to retrain learning models only for small changes to the original problem all point to a potential shortcoming in the classic classification regime, where a comprehensive a priori knowledge of the possible classes or concepts is critical. Without such knowledge, classifiers misjudge the limits of their knowledge and overgeneralization therefore becomes a serious obstacle to consistent performance. In response to these challenges, this paper extends the classic regime by reframing classification instead with the assumption that concepts present in the training set are only a sample of the hypothetical final set of concepts. To bring learning models into this new paradigm, a novel elaboration of standard architectures called the *competitive overcomplete output layer* (COOL) neural network is introduced. Experiments demonstrate the effectiveness of COOL by applying it to fooling, separable concept learning, one-class neural networks, and standard classification benchmarks. The results suggest that, unlike conventional classifiers, the amount of generalization in COOL networks can be tuned to match the problem.

Introduction

The hope in machine learning is often to match or even exceed the performance of humans in tasks once prohibitive to machines. The measure of such success is frequently expressed as the *accuracy* of the learning algorithm on a test set separate from the examples experienced during training. For example, in ILSVRC-2010 (a subset of ImageNet), a dataset with roughly 1.2 million images spanning one thousand of different classes, success is measured on a separate $150{,}000$ images in the test set (Russakovsky et al., 2015). This conventional approach to measuring success then makes it possible to compare machine and human performance if we assess humans on the same test set. In fact, in ImageNet humans achieve about 94.9% on the test set (Russakovsky et al., 2015), a level recently exceeded by deep learning algorithms (He et al., 2016). Thus a narrative emerges whereby machines are creeping past the abilities of humans in visual recognition, a remarkable proposition.

Yet this potential narrative of course oversimplifies the depth of human comprehension hidden beneath the raw accuracy score. Unlike machines, humans *have a sense of the limits of their own knowledge*, a fact not easily exposed through a fixed test set drawn from a similar distribution as the training examples. For example, confronted with adversarially-selected images reminiscent of random noise, a human can easily recognize that the resultant image is not familiar and hence belongs to no known class at all. Deep learning, in contrast, has proven easily fooled by such images into assigning them to a class with over 99% confidence, a phenomenon called fooling (Szegedy et al., 2013; Nguyen, Yosinski, and Clune, 2015).

While being fooled by examples far outside the training distribution raises concerns, such mistakes alone are confined to scenarios that might be dismissed as relatively unlikely. For example, how often in the real world might one encounter an image unlike anything in the real world yet adversarially optimized to trigger a mistaken classification? However, as experiments in this paper will show, fooling is only a foreshadowing of a set of shortcomings in conventional neural networks that fall short of human-level capabilities. For example, a tendency to classify images with confidence that are not from among the training classes means that adding a new class to an already-learned model is unlikely to work, thereby precluding continual learning of new classes. It also means that combining separately-trained classifiers is similarly untenable. Single-class neural networks, which learn to say whether a training instance is within the single class or outside it, are also prohibitive to train for similar reasons.

All these problems result from the tendency of conventional neural networks to draw well-informed discriminative borders between different classes, but to fail at confining the regions of such classes to the span only of the examples in the dataset. The consequences include fooling, but also real practical divergence from the natural ability to learn new classes seamlessly throughout a lifetime. If neural networks could be brought to learn better fitted regions, i.e. ones that do not *overgeneralize*[1], the problems so far cited, from fooling to continual learning, would be largely mitigated or even in some cases solved.

[1] In psychology, in early language acquisition, the same concept is called *overextension* and is often attributed to a limitation in vocabulary (Gershkoff-Stowe, 2002).

To capture this idea this paper later introduces a formal framework for *fitted learning*, which refers to algorithms that learn both good discriminative borders and the distribution of training data at the same time. Moreover, the major contribution of this work is a simple mechanism for achieving fitted learning (and to tune the level of fit) within the conventional apparatus of deep learning and neural networks. The main idea, called the *competitive overcomplete output layer* (COOL), is to assign more than one output to each class and also force such outputs to compete with each other to respond correctly. The resulting dynamic, which is achieved through the usual process of stochastic gradient descent, pushes the competing neurons within a class to respond to overlapping yet non-identical regions of input space, yielding sufficient consensus for confident classification only within a tight-fitting region around the training data. This simple mechanism for fitted learning is thereby a step towards a more natural style of learning in neural networks.

Experiments in this paper in the MNIST digit-classification domain (LeCun et al., 1998) confirm that the COOL approach not only mitigates fooling as expected, but also indeed enables combining multiple separately-learned classifiers into one, and even makes possible effective one-class neural networks, all of which are otherwise prohibitive with conventional convolutional neural networks (CNNs) (LeCun et al., 1998). Moreover, COOL networks actually outperform CNNs with identical architectures (except of course the outputs) in CIFAR-10 and CIFAR-100 (Krizhevsky and Hinton, 2009), not only reaching a higher level of test accuracy, but converging significantly faster (even with the greater number of outputs) thanks to the more opportunistic steps taken by COOL networks through gradient descent.

After covering background material next, the paper introduces the components of the COOL architecture and then intuitively demonstrates their advantage through visualizations of COOL and non-COOL networks in a simple two-dimensional classification problem. The COOL mechanism is then justified more formally and tested within the context of fooling images, after which fitted learning and the connected *generalized classification problem* are also formalized. A series of experiments in MNIST follow on learning separate classifiers that are later combined (which works poorly with conventional CNNs), including the first demonstration of working one-class neural networks. The experiments then conclude with head-to-head classification comparisons in CIFAR-10 and CIFAR-100, followed by a discussion of the implications of fitted learning and COOL.

Related Work

The closest conceptual framework to fitted learning is probably the *open set recognition problem* (OSRP) (Scheirer et al., 2013; Jain, Scheirer, and Boult, 2014) in the vision community, where unknown classes may appear in the testing phase and the learning model should be able to detect such unknowns. However, as a recognition task, in OSRP each learning model is focused on a single class, while the aim in fitted learning is for learning models to handle several classes together. That is, a fitted model should be able to discriminate at the same time as capturing the data distribution for each class. Unlike the common practice in open set recognition, a fitted model should learn a tight halo around different class regions, unless there is a border friction with another class, and therefore should not be very sensitive to the choice of a threshold parameter. Visualization experiments that follow indeed illustrate these properties for COOL in low-dimensional spaces.

An established method that is applied in OSRP and anomaly detection (amongst others) is the one-class SVM (Schölkopf et al., 2001), which tries to capture the data distribution, or more technically the support of the probability distribution, by turning the one-class problem into a classification problem. The main difference between the one-class SVM and fitted models is that the one-class SVM is a discriminative model at its core so it will find a border around some data points rather than capturing their underlying distribution. For example, learning regions in the feature space with a hollow inside is very difficult (if not impossible) for a one-class SVM model.

Fitted learning can also be viewed as inverse generative model learning, where sampling from a particular class involves searching for an input instance that highly activates the corresponding class in the hypothetical fitted model. That is, an input that highly activates an output node in a trained fitted model is an example of the class of that output node and a procedure to generate such instances leads to a generative model. An analogy can be drawn with Generative Adversarial Nets (GANs) (Goodfellow et al., 2014), which try to capture the distribution of data through a generative model that competes against a discriminative model. COOL, as a potential fitted model, induces some form of competition as well, but unlike GANs it does not entail careful training of the competing models because there are no explicit separate models.

Another method with a conceptual connection to COOL is dropout (Srivastava et al., 2014), where an exponential number of models are trained together by randomly dropping out some hidden units during the training phase. Similarly, COOL in effect trains an exponential number of models as well, which also try to learn the same concept. However, COOL also employs a mechanism to encourage diversity in the final subset models and in contrast to dropout it usually results in training speedup.

Approach

COOL in effect expands the traditional output layer by assigning several output units to each class or concept. Furthermore, forcing this group of neurons to compete ensures their diversity by partitioning the input space among output units of the same class. As a result these output units ultimately only agree within regions proximal to the training instances. In other words, unlike conventional architectures, COOL captures the distribution of the training instances of each class *at the same time* as learning to discriminate between the instances of different classes.

The section begins with a description of the basic unit of COOL, the *neuron aggregate*. The remainder of the section

then describes the proposed new architecture for output layers.

Neuron Aggregate

A *neuron aggregate* is a collection of units, called *member units*, that is intended to learn a single concept. In this sense, a neuron aggregate acts as an ensemble and is inspired by the hypothesis (in contrast to the "grandmother cell") that in biological neural networks more than one neuron is involved in recognizing each concept (Gross, 2002; Gazzaniga, 2004).

An *internally-competitive aggregate* is a neuron aggregate whose member units' activations are mutually inhibitory (e.g. through reciprocal inhibition) and is modeled in COOL by a softmax function applied to member units. Note that it accordingly suffices for a neuron aggregate to be part of a softmax layer to become an internally-competitive aggregate.

The Competitive Overcomplete Output Layer (COOL)

A COOL consists of a simple change to the traditional output layer setup of neural networks. In particular, each output unit is replaced by an internally-competitive aggregate. As we elaborate further on this architecture, in the rest of this paper we assume one-hot neural encoding for target output labels. We also define the *expected maximum activation value (EMAV)* of a unit as the expected[2] highest activation value a single unit can have; for example, normally a traditional logistic sigmoid saturates at an EMAV of one. An output unit in a softmax layer normally is trained with values of either 0 or 1, indicating the deactivation/activation of the corresponding unit, respectively. In our framework, because a traditional unit is replaced by a neuron aggregate, COOL accepts lower values to indicate the activation of member units. More specifically, $1/\omega$ indicates activation of a member unit and zero indicates its deactivation, where ω is a hyperparameter called the *degree of overcompleteness (DOO)* that indicates the number of member units in a neuron aggregate. In this paper all neuron aggregates of the same output layer use the same degree of overcompleteness. In effect, the EMAV value of each member unit of an aggregate is $1/\omega$.

During the training phase, all the member units of the same neuron aggregate are trained with the same value, i.e. zero or $1/\omega$, depending on the desired activation of the corresponding aggregate. Algorithm 1 shows this procedure formally. Note that $1/\omega$ combined with the one-hot encoding of class labels ensure that the sum of all the member units always equals 1.0; a cross-entropy cost function is then applied straightforwardly to train such a network. The activation values of member units in a COOL setup can be interpreted as the scaled probability of inclusion of the input instance in a particular class.

It is important to highlight the role of the two components of a COOL layer, i.e. overcompleteness and competitive training through the softmax function. The first component

[2]The word "expected" conveys that maximum activation value a unit can take during training may differ from the maximum value it can output during testing.

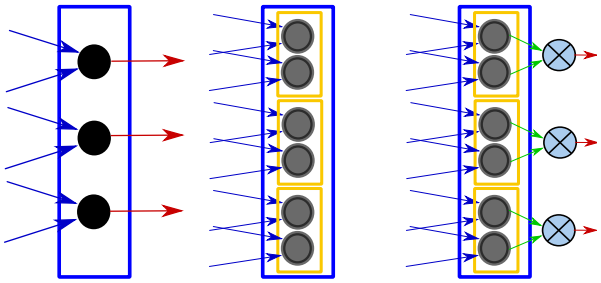

(a) Conventional output layer (b) COOL during training (c) COOL during testing

Figure 1: **The Output layer in a conventional neural network vs. COOL.** (a) A conventional output layer has three neurons to learn three concepts. (b) An overcomplete output layer for the same task is shown during training. Each yellow rectangle depicts an aggregate and two gray units in each aggregate show that all aggregates have DOO of 2 in this architecture. (c) The same COOL output layer computes its decisions through multiplication during the test phase.

in effect adds new dimensions to the output layer while the second enforces a continuing competition among the member units of the active neuron aggregate. The section on *Why the COOL Mechanism Works* shows in detail how these two components enforce a beneficial competition among member units. Visualization experiments that come first provide initial intuition on how this competition can lead to a fundamental change in output layer behavior, in effect preventing output units from overgeneralizing into regions far away from the training instances. These experiments will suggest that internally-competitive aggregates are able to learn an appropriate probability distribution of their assigned concepts, a critical property for many learning tasks.

During the inference phase, the outputs of the member units of an aggregate are combined through *multiplication* to calculate the total activation value of the aggregate. The idea is that member units are trained to minimize variance and therefore are expected to exhibit the least variance and highest activation when the network inputs are within the proximity of training instances. At the same time, softmax enforces a constant total sum of the member units. Figure 1 compares a conventional output layer with a COOL, including the multiplications that happen for inference.

Multiplication helps to maximize activation (and minimize variance) of member units at low computational cost because the product of a set of variables constrained by a finite sum is maximized when the variables are all equal. Finally, to control the degree of generalization of a particular class c, the output values can be exponentiated by a softness parameter σ_c after application of the multiplication. This operation allows tuning the decision boundaries of individual classes arbitrarily. The role of this parameter will be more clear when separable concept learning is later revealed as a consequence of COOL. Optionally, to get final results in the form of probability values each aggregate output can be multiplied by ω^{ω}. Algorithm 2 describes the procedure for

Algorithm 1: How the target for a training instance is set in COOL.

Input: degree of overcompleteness ω, class index for a training instance η, total number of classes k
Output: training target l
$l \leftarrow \text{zeros}(\omega \times k)$
for $i \leftarrow 1$ ***to*** ω **do**
  $l[(i-1) \times k + \eta + 1] \leftarrow \frac{1}{\omega}$

Algorithm 2: Inferring a test instance in COOL.

Input: degree of overcompleteness ω, prediction vector p, softness parameter vector σ, number of classes k
Output: test label probability vector l
$l \leftarrow ones(k)$
for $i \leftarrow 1$ ***to*** ω **do**
  for $j \leftarrow 1$ ***to*** k **do**
    $l[j] \leftarrow l[j] \times p[j + (i-1) \times k] \times \omega$
for $j \leftarrow 1$ ***to*** k **do**
  $l[j] \leftarrow l[j]^{\sigma_j}$

calculating the activation value of an aggregate during testing.

Visualization Experiments

In this section a low-dimensional artificial dataset helps to visualize the behavior of COOL networks and compare them with traditional neural networks. These experiments suggest that COOL networks are able to successfully capture the underlying distribution of training instances while still preserving the generalization ability of conventional neural networks. In other words, *they construct a notion of the limits of their knowledge*; they can recognize instances that are not inside their scope of expertise (i.e. far from the training points) in addition to behaving like a traditional learning model.

Problem Definition

A major obstacle to understanding neural networks is the inability to visualize their behavior over a large cross-section of the input space. To address this challenge, a simple two-dimensional classification problem is introduced that allows such broad visualization. The problem involves classifying points that are located only within the perimeters of two concentric circles of points in the instance space, c_1 and c_2, with different radii and both centered at $(0, 0)$. The training set includes only points within these two bands (dashed circles in figure 2). Each point within circle c_1 (inner circle) is a member of class one and each point within circle c_2 (outer circle) is a member of class two. A uniform sampling of these circles generates the training set.

Visualization experiments in this section train neural networks with different architectures on this two-class problem. The trained network's output for all possible two-dimensional instances at a high resolution is then plotted on a grid. That way, it is possible to see how the network classifies instances outside the training distribution.

Traditional neural networks and other discriminative classification models look for a discriminative border. For a two-class problem this objective means that they will typically partition the whole input space into two regions, both assigned to one of the available classes. One possible choice for such partitioning is a max-margin border, which is depicted in figure 2 as a red line.

In the first experiment a shallow architecture (one hidden layer of 200 neurons) is compared with a deep one, consisting of three hidden layers of 400, 300, and 200 neurons, respectively (neither with COOL). Figure 3 depicts the activation map of the output unit assigned to the inner circle in both the shallow and deep neural networks. These pictures show how adding more hidden layers result in finding a better decision boundary. However, the deep approach also tends to preserve high activation of output units in the regions that are far from the training points, a problem that can be described as *overgeneralization*. In classification problems where every possible input instance is a member of the set of known classes overgeneralization does not exist by definition, but in most real-world applications:

- the complete set of classes is *not* known and
- not all the possible inputs are an instance of *any* class.

In fact, *most* possible inputs may not be valid and only those input instances that are likely to be drawn from the training data distribution should be considered valid. One naive solution to the overgeneralization problem is to add an extra class including those inputs that are not a member of any class. There are two major drawbacks with this approach:

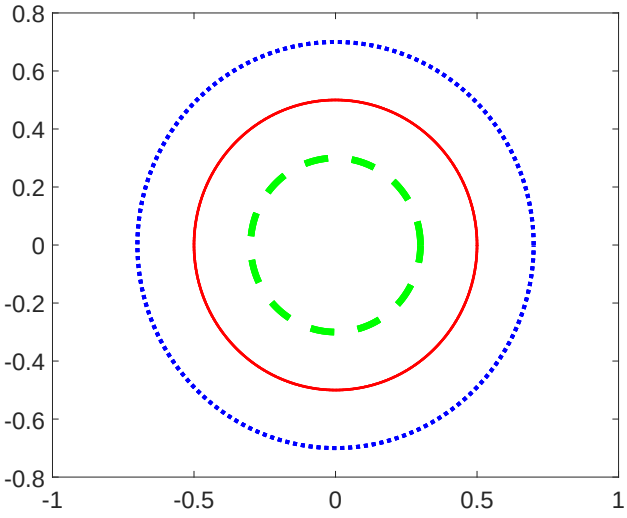

Figure 2: **Illustration of the training data for the visualization problem.** The blue and green dots show the instances from two different classes and the red circle depicts the max-margin border.

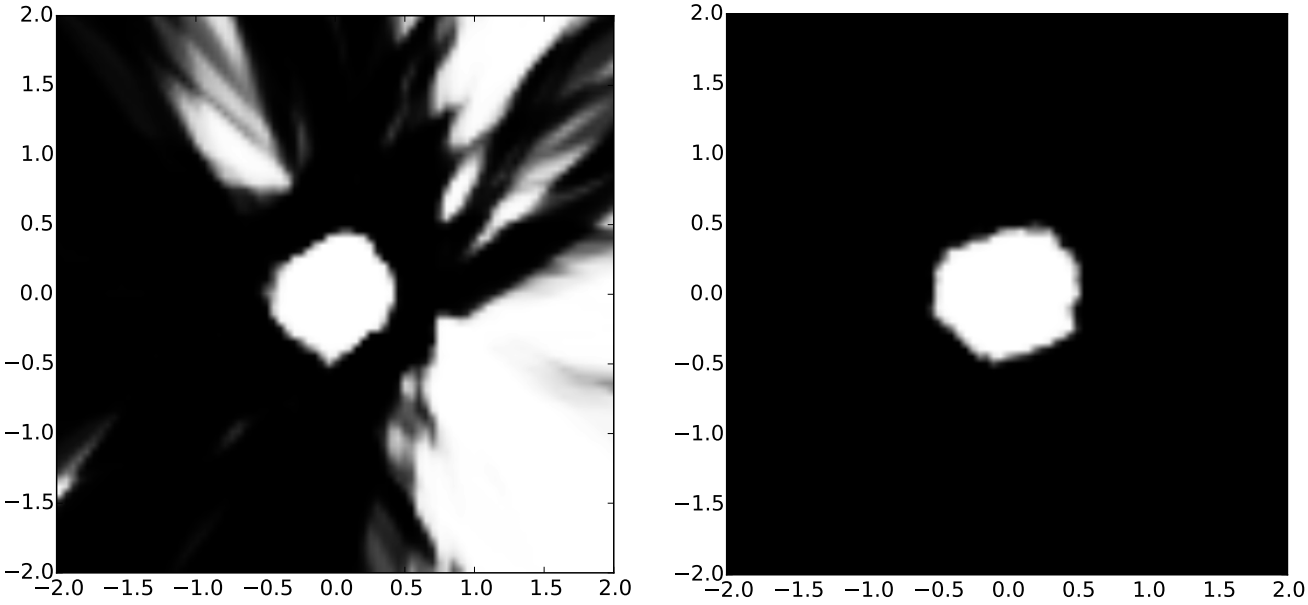

(a) Shallow network

(b) Deep network

Figure 3: **Generalization ability of conventional deep vs. shallow neural networks on the visualization problem.** The activation map of the output unit assigned to learn the points within the inner circle is shown for both networks. (The activation map of the outer circle, which is not shown, is the complement of these plots.) The highly-activated area inside the circle in both plots indicates the ineffectiveness of conventional neural networks in preventing overgeneralization. (a) The abundance of highly activated points outside the inner circle is usually interpreted as a shortcoming of shallow neural networks in generalizing well. (b) The activation map of the deep architecture finds a better approximation of the max-margin decision border. While this outcome is conventionally considered a good behavior, it also hints that deep models are easier to be fooled.

- Accurate generation of such instances necessitates access to the underlying training instance generation process.
- Usually this set of instances has a much larger size, which can result in a dramatic data imbalance.

In contrast, COOL provides an alternative solution that does not require any modification to the training data. To demonstrate this capability, in the next visualization a COOL network is trained on the same two-circle problem. Figure 4 depicts the activation maps of the two output units corresponding to inner and outer circles, respectively (for both a COOL network and conventional network for comparison). These pictures clearly show how the COOL mechanism can effectively prevent overgenaralization, suggesting the ability of these types of network to capture an *implicit* understanding of the data generation process.

These experiments so far may raise the concern of whether COOL networks are too restrictive to generalize well. Accordingly, the next set of experiments suggest three means to modulate generalization in COOL, namely:

- changing the degree of overcompleteness,
- changing the depth of the architecture, and
- changing the softness parameter.

In short, by choosing the right hyperparameters one can tune the generalization ability of a model from pure discriminative networks on one end of the spectrum to behaving like a histogram approximation of the data distribution on the other side. This insight paves the way toward *fitted learning*, i.e. learning models with the right amount of generalization.

Figure 5 depicts the activation maps of the same concept in two trained networks with different degrees of overcompleteness (*DOO*) on the two-circle problem. Both networks have the same architecture except for the output layer. These visualizations suggest that higher DOO can result in less overgeneralization. In other words, the less the DOO the more the network is likely to generalize over unseen regions. This observation is also supported by the intuition that a DOO of one is simply a traditional neural network.

Next, two COOL networks with different depth but the same output layer (with DOO $= 5$) are compared. Figure 6 shows how adding more hidden layers can lead to more generalization ability but a higher risk of overgeneralization. In fact, in this experiment sometimes deep architectures lead to generalization over unwanted regions (e.g. partially filling the circle shown in figure 6).

Finally, a brief experiment with the Iris dataset (Fisher, 1936) provides the opportunity to show the behavior of COOL in the presence of *sparse* datasets. In general, sparse datasets, as opposed to *dense* (or redundant) datasets, are

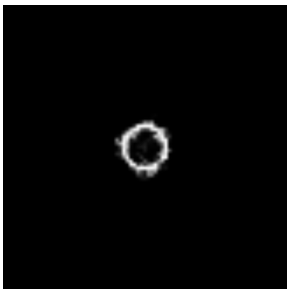

(a) COOL inner circle output

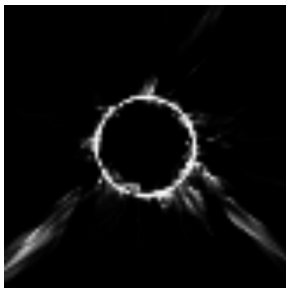

(b) COOL outer circle output

(c) Conventional neural network inner circle output

(d) Conventional neural network outer circle output

Figure 4: **The activation map of output units in a COOL vs. a conventional neural network (hollow circle is ideal).** Images (a) and (c) depict the highly active area for the inner circle in a COOL versus a conventional MLP. Notice the presence of the donut hole in (a). Images (b) and (d) show the highly active region for the outer circle. Again the highest outputs for COOL are within the correct band around training data while the conventional neural network generalizes to infinity. These plots depict how COOL can prevent overgeneralization at the same time as learning to solve the problem.

those with few training instances[3] such that even excluding a few examples can lead to significant negative impact on training. While deep learning techniques are mainly effective on dense datasets, sparse datasets are still important because in many domains training instances are scarce.

For visualization purposes, in this experiment only the first two input features of the Iris dataset are included in the training set. Figure 7 depicts the dataset in this reduced space of features whereas the activation maps of output units of COOL networks assigned to three different types of iris are shown in figure 8. These plots show how changing softness parameter can shrink/extend the activation regions of each class, which can lead to a significantly better generalization.

In conclusion, the experiments in this section provide intuitive empirical evidence of the ability of COOL to learn the underlying data distribution of training instances while retaining the usual discriminative ability.

[3]To be more precise, the density of training instances is low.

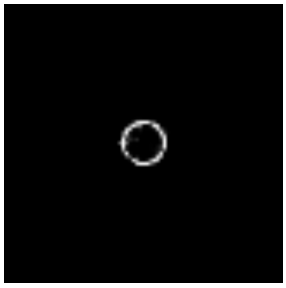

(a) DOO $= 10$

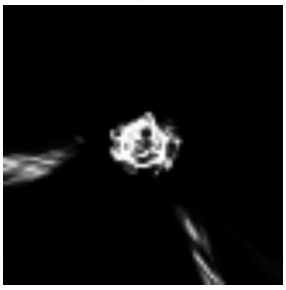

(b) DOO $= 2$

Figure 5: **Effect of changing DOO on the generalization ability of COOL.** Both plots show the activation map of the aggregate assigned to learn the inner circle.

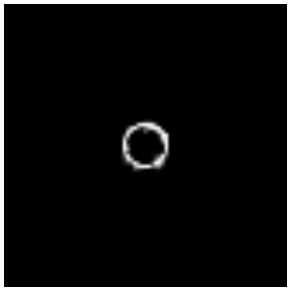

(a) 1-layer COOL

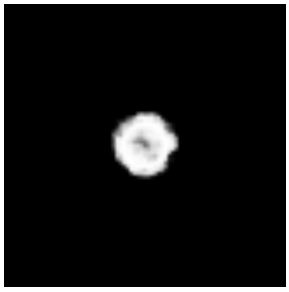

(b) 5-layer COOL

Figure 6: **Effect of applying deeper architectures on the generalization ability of COOL.** The activation maps associated with the inner circle in COOL networks with 1 (a) and 5 (b) hidden layers are shown, respectively.

Why the COOL Mechanism Works

This section explains how the two components of the COOL setup, i.e. overcompleteness and competition, can work together against overgeneralization. In particular, it turns out that the overcomplete architecture leads to simultaneous training of an exponential number of models together while competition preserves diversity in the member units of a neuron aggregate that are trying to learn the same concept. The argument begins with an aggregate X with n member units, $X = \{x_1, x_2, ..., x_n\}$.

Overcompleteness

The idea behind the overcomplete layer is that it produces a dynamic akin to training an exponential number of models all at once, which echoes the motivation behind dropout (Srivastava et al., 2014). Assume that the output of each member unit is scaled just before multiplication in the test phase, i.e. the output of member unit x_i $\forall i \in \{1, 2, ..., n\}$ is transformed according to $S(x)$ =max$(\omega \times x, 1)$, where ω is the DOO.

In this realization each member unit's output can be interpreted as a probability value. Furthermore, any subset of an aggregate, except for the empty set, could in principle

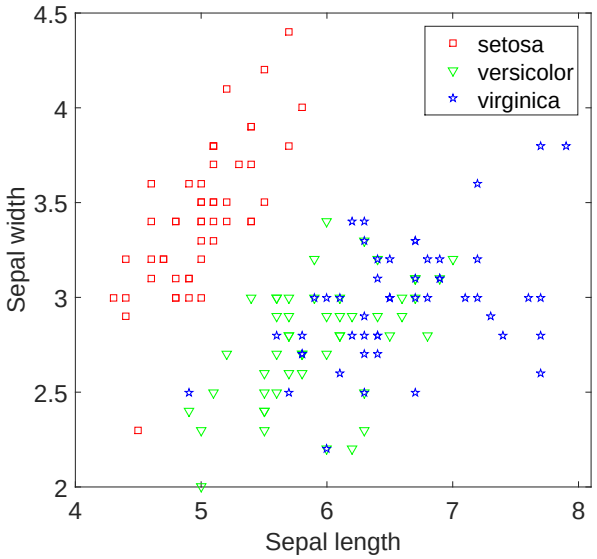

Figure 7: **Iris dataset as an example of sparse dataset.** Note that the reduced input dimensionality in this problem makes it much harder than the original Iris problem.

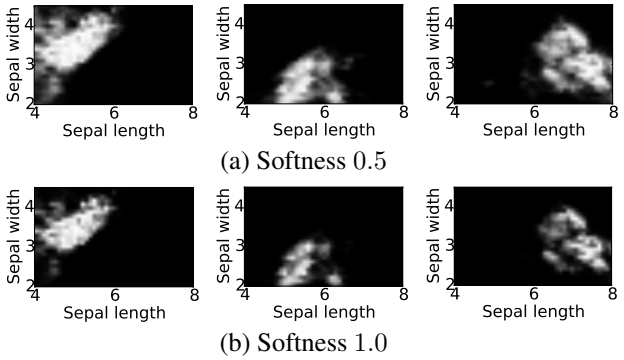

(a) Softness 0.5

(b) Softness 1.0

Figure 8: **How generalization ability can be tailored in COOL by changing the softness parameter.** The top three pictures (a) depict the activation map of three different classes in a COOL network when the softness parameter is 0.5. The bottom pictures (b) depict the activation maps in the same network when the softness parameter is 1.0.

replace the original aggregate as the decision-maker. Consequently, training an aggregate is hard because it is analogous to training of $2^n - 1$ neuron sub-aggregates with the same activation behavior at the same time. This challenge is further exacerbated by having an aggregate for each class in COOL.

Competition

Competition among the member units of the same aggregate is the instrumental mechanism of COOL that yields its protection against overgeneralization. Recall that during the training of a particular instance there is only one *active aggregate*, i.e. the aggregate whose member units are trained for non-zero values. Because all the member units in the active aggregate are trained to learn the same function, one might expect them to converge to the same weights for each neuron in the aggregate. Interestingly, the competition induced by softmax prevents such an outcome.

To elaborate, assume X is the current active aggregate for the input instance, χ; X has different member units and in practice in gradient descent there is at least one non-zero gradient[4] amongst $x_i \ \forall i \in \{1, 2, ..., n\}$. The softmax function then implies that within a neighborhood N of χ, $\xi \in N(\chi)$,

$$x_1(\xi) + x_2(\xi) + ... + x_n(\xi) \simeq 1, \tag{1}$$

where $x_i(z)$ is the activation of x_i for the input instance z. Note that if all the member units are similar to each other, it is likely that they are all within a close proximity in the space of the cost function. This closeness stems from all member units applying the same cost function, having the same input, and targeting the same value. Therefore, in a neighborhood N of χ equation 1 yields:

$$\frac{\partial E}{\partial x_n} \simeq -\frac{\partial E}{\partial x_1} - \frac{\partial E}{\partial x_2} - ... - \frac{\partial E}{\partial x_{n-1}}, \tag{2}$$

where E is the cost function. In other words, the gradient of x_n is *competing* (i.e. pulling in a different direction) with the gradients of all the other member units. Ignoring the specific parameters to each member unit, there are two possible cases:

- The final parameter updates follow gradient descent for all the neurons $x_1, ..., x_{n-1}$, which results in a gradient ascent for neuron x_n.
- There is a member neuron in $\{x_1, ..., x_{n-1}\}$ whose gradient descent fails in the new parameter updates.

For these reasons, converging to the same unit is not attractive to the member units when starting from different initial parameters.

In short, assigning several competing units to the same concept can be thought of as a *game of changing model parameters* between member units of the active aggregate that can lead to different partitioning of the input space among them. As a result we expect these units to preserve their diversity for arbitrary inputs far from the training instances and only to maintain consensus over the regions proximal to the training set. To further illustrate this principle, figure 9 depicts the activation map of a COOL network that lacks the competition component. This neural network is trained on the two-circle problem from the visualization experiments. Compared to COOL (figure 4), figure 9 shows that without competition all the member units associated with the same concept can easily converge (and hence overgeneralize). Finally, note that the two plots in figure 9 are not the exact complement of each other, which indicates that even without competition there is still some advantage over conventional neural networks.

[4] By "gradient" we mean the gradient of cost function with respect to the member units.

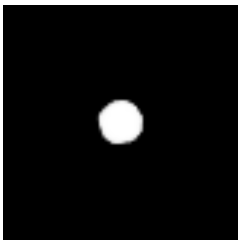

(a) Inner circle output

(b) Outer circle output

Figure 9: **The effect of removing competition from COOL.** These plots depict the activation map of aggregates corresponding to (a) inner and (b) outer circles in the two-circle problem when softmax is removed (compare to figure 4).

The Fooling Problem

A phenomenon called *fooling* has recently attracted attention in the deep learning community (Szegedy et al., 2013; Nguyen, Yosinski, and Clune, 2015). The main observation is that neural networks can be tricked into outputting with confidence the wrong class for instances well outside their training distribution. In other words, the fooling problem is a direct consequence of overgeneralization, which is more prominent in deep architectures.

To clarify the fooling problem we begin with a simple example: The classification task is to determine whether a fruit is a watermelon or an apple based on the weight of the fruit in kilograms. A possible training set is then $T = \{(0.25, apple), (0.2, apple), (5.25, watermelon), (7.3, watermelon)\}$. In this example a discriminative model, such as a SVM, will find a decision boundary that can be translated into $if\ \ weight > 2.75\ \ then\ \ watermellon,\ \ otherwise\ \ apple$. Often discriminative models can output probability values based on the underlying decision boundary as well (Duda, Hart, and Stork, 2012). For example, in such a model if an input is close to 2.75, say 2.5, it would be classified as an apple but with a lower probability than a fruit with weight 1.5. The fooling problem occurs when the input is not drawn from the underlying data distribution, e.g. if we try a fruit that weighs 50 kg most discriminative models trained as above would classify it as a watermelon even though such a fruit possibly does not exist.

It is worth noting that ensemble techniques on their own have almost no mitigating effect on the fooling problem (Szegedy et al., 2013) because they merely generate a better decision boundary, in the form of a larger margin. However, fooling is not a problem of the margin.

The biggest downside of fooling is the fact that it is *dataset-dependent* and not model-dependent. In other words, one can easily find fooling examples that can fool almost *all* discriminative models only based on a limited knowledge of the training set (Szegedy et al., 2013; Nguyen, Yosinski, and Clune, 2015).

This section first reviews reasons behind fooling. Then it provides a simple way to measure how easily a trained neural network can be fooled based on a generative adversarial approach. Finally, it compares the fooling in COOL versus conventional neural networks on the MNIST dataset. The results suggest that the COOL mechanism is capable of preventing fooling in neural networks.

Reasons Behind Fooling

Goodfellow, Shlens, and Szegedy (2015) postulate a linear explanation for the existence of adversarial examples, where small perturbations to an existing training example can lead to significant changes in the output in a high-dimensional problem. They also extend this explanation in their appendix to the more general case of an arbitrary fooling example, i.e. one not necessarily similar to a training example, such as in Nguyen, Yosinski, and Clune (2015). This paper (and COOL) focuses on the latter case.

In general, suppose that training instances are sampled from an unknown probability distribution D. Consider the classification of an arbitrary input example x by neural network N; x is a fooling example if it is classified with a high confidence by the network N while *not* being drawn from the probability distribution D. A significant factor behind fooling is then that discriminative classifiers tend to overgeneralize. That is, they classify with confidence many input instances regardless of their resemblance to the training data. Several active output units for a single instance implies low confidence of the classifier, but a single dominant output unit for an outlier example implies fooling the model. Because these kinds of learning models implicitly draw borders in the instance space, the distance between input points in the same region is in effect ignored, resulting in the same degree of classification confidence for input instances that can be far away. Interestingly, COOL networks inherently address this issue.

Generating Fooling Instances

Previous methods for generating fooling images from Szegedy et al. (2013) and Nguyen, Yosinski, and Clune (2015) search directly for images that fool the network. In contrast, here we introduce a third option that does not require solving a constrained optimization problem: A random input instance x is fed into a new trainable neural network g, called the *fooling generator network* (FGN), whose output is passed to the actual model f. In other words, $g(x)$ is the fooling input to the network instead of x. Gradient descent can train network g such that $g(x)$ generates a good fooling image.

During this procedure, x and parameters of f are fixed and only g is being trained. More formally, $f : \Re^n \rightarrow [0,1]^k$ is a mapping from input vectors to a probabilistic target vector space ρ, $h : \rho \rightarrow \{1, ..., k\}$ is an invertible mapping from ρ to a discrete set of labels, and $g : \Re^m \rightarrow \Re^n$ is a mapping from input vectors of size m to vectors of size n (note that m can be equal to n). The key idea is thus to solve the optimization problem $Minimize||f(g(z)) - h^{-1}(l)||_2$, for a given input $z \in \Re^m$ and target label $l \in \{1, ..., k\}$, where $h^{-1}(.)$ denotes the inverse of $h(.)$. Note that $||.||_2$ can be re-

placed by cross-entropy in this optimization problem, which is the procedure in the experiments that follow.

This approach has several advantages: First, there is no need for *constrained* optimization because one can control the outputs of network g by choosing the right activation function. For example, a sigmoid function can give an output in the range $(0, 1)$. Second, a good choice of architecture can indirectly impose desirable constraints on the generated fooling examples, e.g. when dealing with images, a convolutional FGN g imposes some natural image properties on the generated fooling images. Finally, almost all the components of this approach are already provided in most machine learning packages.

MNIST Fooling Experiment

In this experiment first a COOL network and a conventional CNN are trained on the MNIST dataset. Then 20 trials are attempted to generate fooling instances for each model. A trial consists of training a FGN activated with a random input to trick the model and is considered successful if the model classifies the generated fooling image with more than 99% confidence. However, if such a fooling example is not found before $10{,}000$ parameter updates of the FGN then the trial is a failure. More details on this experiment are in the Appendix.

COOL preserves the generalization ability of CNNs for this task: the classification accuracies of COOL and conventional CNNs were 99.18% and 99.14%, respectively. At the same time, overall, in 200 trials (20 for each digit), the FGN approach fooled the conventional CNN with a 100% success rate versus a 47% fooling rate for the COOL CNN. Figure 10 depicts examples of typical generated fooling images for both the conventional and the COOL CNN. Interestingly, on average a conventional CNN is fooled before 13 updates of the FGN parameters whereas the COOL CNN needs more than $5{,}000$ updates, which suggests a drastic shrinkage of the high confidence classification regions within the COOL CNN.

Another interesting observation is that most often (more than 60% of the time) the conventional CNN confidently classifies a random image (uniform random noise without any modification) as digit 8 while COOL is *never* fooled by a random image, which is further evidence of the severity of the fooling problem in conventional neural networks. Also, while the conventional CNN has a 100% fooling rate for each individual digit, the FGN approach exhibits variable success for different digits with COOL. For example it failed to trick the COOL CNN for the digit 1 in all of its attempts. This result suggests a more effective approach to preventing fooling may involve a specific DOO for each aggregate.

Finally, it is important to note that this experiment assumes complete access to the models that are meant to be fooled. In a real-world scenario, fooling is likely based on limited access to the training set or some input/output samples from the trained model. The low fooling rate of COOL networks in this experiment with complete access to the model therefore hints that they may in effect prevent fooling in many real scenarios.

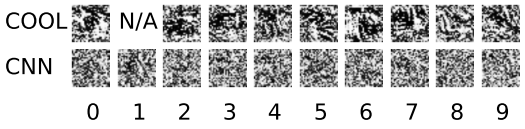

Figure 10: **Example fooling images generated by the fooling generator network for the CNN.** Left to right, each image represents digits 0 to 9, respectively. These images are generated by a FGN consisting of a fully connected layer to produce plausible images. They are classified by the trained conventional CNN (bottom) and COOL CNN (top) with at least 99% confidence. The top image for 1 is blank because the FGN could not generate any fooling image in all of its trials for this digit. Even though none of these images are close to real digits a qualitative difference between the images on top and bottom is evident.

Fitted Learning and Separable Concept Learning

So far COOL has been offered as a mechanism to prevent overgeneralization in neural networks. As a consequence COOL networks are able to prevent fooling or at least make it relatively difficult for an adversary. A natural question that follows is how to measure the amount of overgeneralization in a model. This section defines such a measure by introducing the *generalized classification problem* that leads to a notion of *extended empirical risk*, which is a natural extension of empirical risk (Vapnik, 2013).

Recall from classical statistical learning theory (Vapnik, 2013), where one assume access to a finite collection of i.i.d. samples in a training set, $T = \{(x_i, y_i)\}_{i=1}^{n} \sim D(X, Y)$, where X is the input space of training set T, Y is the output space, and D is the data distribution over $X \times Y$. In the canonical classification problem, we also assume Y is a finite set of objects of size c.

Given a non-negative, real-valued loss function $L(y', y)$ that measures the loss of predicting y as y', the risk R of a function $f : X \to Y$, called a *hypothesis*, is defined as:

$$R(f) = E(L(f(x), y)). \tag{3}$$

The goal of classification is to find a hypothesis f^* such that $R(f^*)$ is minimal. Note that every hypothesis can be associated with an underlying interpretation of $P(Y|X)$. To avoid trivial solutions, hypotheses are usually restricted to a class of functions F. In other words:

$$f^* = \text{argmin}_{f \in F} R(f). \tag{4}$$

Solving equation 4 is theoretically impossible because of ignorance about $D(X, Y)$; thus an approximation of $R(f)$ is expressed by $R_e(f)$:

$$R_e(f) = \frac{1}{n} \sum_{i=1}^{n} L(f(x_i), y_i). \tag{5}$$

Finding f^* using the R_e instead of R is called *empirical risk minimization*, which is a bedrock principle of statistical learning theory (Vapnik, 2013).

Introducing the Generalized Classification Problem

In conventional classification, a learning model simply draws a border between different concepts. One drawback of this strategy, as discussed earlier, is an inherent overgeneralization. *Generalized classification*, introduced here, addresses this issue by asking for the right amount of generalization. Moreover, in many real-world applications a comprehensive knowledge of the concepts is not available during training, or even worse, some new concepts might emerge over time. For example, in classification of a certain fish species we cannot rely on our current knowledge of fish because new species are discovered over time. Thus, ideally we want our models to adapt to new emerging concepts without significant retraining.

Another aspect of conventional classification is its intuitive interpretation. It is a fair assumption that classifiers are more confident when dealing with familiar test instances whereas their confidence degrade amid novel cases. However, counter-intuitively, learning models' confidence is merely based on the test instance's distance to the decision border and its similarity to training instances has insignificant effect on it. The idea of generalized classification is to remedy this limitation by taking into account the underlying training data distribution when outputing confidence values.

Finally, arguably the final role of learning models should not end with their assigned classification task but rather they should be components of a broader construct. Generalized classification aims to provide models that perform subtasks while not conflicting with other components of the broader construct and also properly responding to unexpected circumstances.

In the canonical classification problem we assume a predefined known set of objects Y that correspond to the class labels available in the training set. In contrast, in generalized classification we assume access only to a subset of the objects in Y, $\Theta \subseteq Y$, in the training set while the actual size of Y is unknown. This problem is closely related to the *open set recognition problem* (Scheirer et al., 2013), where a learning model can respond to unknown classes during the test time. However, unlike open set recognition, generalized classification involves learning an arbitrary number of concepts simultaneously by a single learning model that plays the dual roles of density estimation and discrimination. Similarly to the definition of risk we define *extended risk* as follows:

Given a non-negative real-valued loss function $L(y', y)$ that measures the loss of predicting y as y' and a function $g : X \rightarrow \{\Theta \cup \varepsilon\}$ that we call a *strong hypothesis*, the extended risk R^{ε} of g is defined as:

$$R^{\varepsilon}(g) = E(L(g(x), y)), \tag{6}$$

where $\varepsilon = Y - \Theta$ and thus $D(x, y = \varepsilon) \equiv D(x, y \notin \Theta)$. In other words, a strong hypothesis is a conventional hypothesis when x is an example of a known concept to the model, and a constant mapping $g(x) = \varepsilon$ otherwise. Informally, ε is an augmented member to the set of known classes that represents all the unknown concepts and a strong hypothesis may assign a test instance to this set by outputting very low probability of inclusion to all known classes. This approach contrasts with a conventional hypothesis that assigns any test instance to one of the known classes without an unknown option.

Generalized Classification and Fooling

Deactivation of output units for points that are not in the proximity of training data in the initial experiments in the two-dimensional space of the two-circle domain suggested that COOL networks are able to provide strong hypotheses. Further experiments then showed that they can significantly reduce fooling. It is worth highlighting the relationship between strong hypotheses and fooling. In particular, to fool the learning model g one should find an instance $x \sim D(x|y \notin \Theta)$ such that $g(x) \in \Theta$. However, by definition, if g is a strong hypothesis with low extended risk then it tends to assign such an x to ε. Therefore, the closer a learning model is to a strong hypothesis with low extended risk the less likely it is to be fooled. Note that in this discussion we distinguish *adversarial examples*, which are drawn from $D(x|y \in \Theta)$ but not classified correctly, from arbitrary fooling examples. Also note that the same proposition is not valid for a hypothesis that is not strong because:

- by definition it has to assign x to a member of Θ and
- there is no obligation imposed on the output of a hypothesis by the loss function when $x \sim D(x|y \notin \Theta)$.

Separable Concept Learning

Unfortunately it is not straightforward to formulate the generalized classification problem using an approximation of extended risk through samples in the training set because we have only partial knowledge about Y. However, it is possible to evaluate the performance of a strong hypothesis in the framework of classical statistical learning. This evaluation consists of two components: (1) model performance on instance x when $x \sim D(x|y \in \Theta)$, and (2) when $x \sim D(x|y = \varepsilon)$; we call the latter the *inhibition* ability of the model because it must inhibit its tendency to activate at least one output class. While the first component can be evaluated by any classification metric, e.g. classification accuracy, for the inhibition ability we introduce the idea of *separable concept learning* (SCL). The advantage of this new metric over previous evaluation procedures applied in open set recognition and object detection tasks (Scheirer et al., 2013) is that it does not rely on any threshold and measures the rejection ability of the model implicitly, which simplifies the evaluation procedure and reduces the processing time significantly.

Intuitively, a strong hypothesis should be able to learn different concepts separately from each other because it is aware of the limits of its own knowledge. As a result, these models can learn new nonoverlapping concepts simply by accepting new output nodes from a separately trained strong hypothesis. In other words, it is possible to train several strong hypotheses over different concepts and merge them together to get a single model over the whole set of concepts. This capability in turn makes *knowledge transfer* possible. Imagine a situation where one has two models to rec-

ognize, {people, animals, plants} and {sea, ground, sky}, respectively. Then a model that finds people in the sea can be simply built by applying a conjunctive operator to the corresponding output nodes in the two models. In a similar manner disjunctive concepts (and concept complements) can be built based on previously learned concepts. Another closely related subject is *incremental concept learning* (ICL), where concepts are presented to the learning model incrementally. Separable concept learning can be conceived as the extreme case of ICL where after learning a few concepts their training instances diminish from the memory. Interestingly, ICL and SCL are more reminiscent of human learning than the current full concept learning approach. Furthermore, in most realistic applications the size of Y is unknown at the beginning, or could grow over time, which makes ICL a practically important subject in supervised learning.

Formally, to evaluate the inhibition ability of a model, let $concat$ be the concatenation operator and ρ a partition of Θ such that every subset in ρ is at least of size 2 (contains at least two concepts from Θ). The *extended empirical risk* of a strong hypothesis g with respect to ρ, $R_e^\rho(g)$, is then defined as:

$$R_e^\rho(g) = \frac{1}{n}\sum_{i=1}^{n} L(\text{concat}_{\forall k \in \rho}(g_k(x_i)), y_i), \qquad (7)$$

where we assume each strong hypothesis g_k is a learning model trained on the subset of training set samples whose classes are included in the corresponding ρ. Note that in case of $\rho = \{Y\}$, equation. 7 will reduce to the classification accuracy measure. In addition, merging outputs of different models usually necessitates having outputs in the form of probability values.

Abstractly, while comparing two learning models, L_1 is strictly better than L_2 if $\forall \rho R_e^\rho(L_1) \geq R_e^\rho(L_2)$. However, the possible ways of partitioning Θ grows exponentially with the size of Θ. Thus in practice an approximation is used by sampling from all the possible partitions.

Informally, to apply equation. 7, one can train several components of a SCL model on different subsets of the set of classes and measure the performance of the whole (merged) model over the test set. This approach provides a simple procedure to evaluate different learning models in terms of acting as a strong hypothesis.

MNIST Separable Concept Learning Experiment

To evaluate the feasibility of SCL and consolidate the results from the fooling section, a number of new experiments are conducted with the MNIST dataset. The results show that the COOL mechanism can dramatically improve the ability of neural networks to learn subsets of concepts separately. These results further support that COOL networks are able to successfully prevent overgeneralization in high-dimensional spaces.

Experimental Setup

Two CNN architectures are compared that are identical other than their output layers. Five instances of each variant are then trained on different subsets of MNIST. More specifically, we partition the training set T into five subsets, T_1, T_2, T_3, T_4, and T_5, where $T_1 = \{(x, y) \in T \mid y \in \{0, 1\}\}$, $T_2 = \{(x, y) \in T \mid y \in \{2, 3\}\}$, etc. Then we train a CNN from each architecture on T_1, one on T_2, and so forth. In other words, $\rho = \{T_1, T_2, T_3, T_4, T_5\}$ in equation 7. More training details can be found in the Appendix.

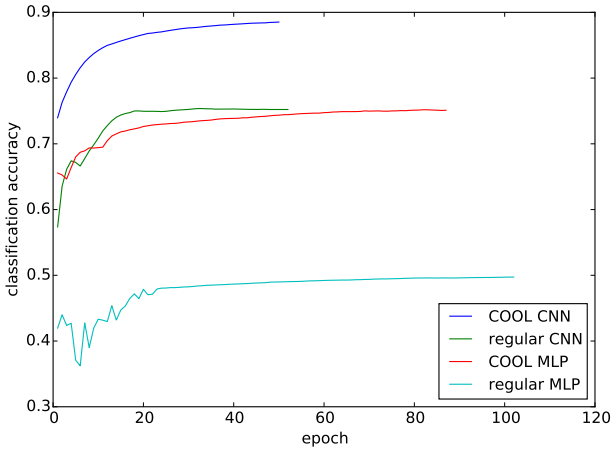

Figure 11: **Test performance of different learning models in separable concept learning in MNIST.** In this experiment the performance of five neural networks, each trained on a different pair of digits from MNIST, is evaluated when they are merged. The main result is the improvement of COOL over conventional neural networks.

Results

Figure 11 shows the average test/validation error (equation 7) of the conventional CNN versus COOL CNN over epochs. Here the COOL models significantly outperform conventional CNNs with about a 13% improvement in recognition rate (over ten runs). Figure 11 also shows the accuracy of MLP architectures (over two runs) on the same task, where an improvement in performance of about 25% is achieved simply by applying the COOL. Interestingly, the COOL MLP (without convolution) achieves competitive performance with the conventional CNN on this task.

Overall, these results suggest that COOL can significantly improve the inhibition ability of learning models, which in turn leads to a more accurate representation of the knowledge embedded in the dataset and robustness of the learned concepts.

One-class Neural Network

Based on the COOL architecture and following a similar approach to a one-class SVM, this section proposes one-class neural networks, yet another possibility created by COOL. This alternative one-class learning model enables direct application of deep learning methods in the field of one-class recognition. The experiments below suggest that this approach is promising and can allow neural networks to enter

the problem domain of estimating the support of probability distributions.

Approach

To capture the underlying distribution of a set of unlabeled data, Schölkopf et al. (2001) proposed the one-class SVM, where the origin is treated as the *only* member of the second class. The algorithm then tries to find a maximum margin between the mapped features of the instances of the first class and the origin. To the best of our knowledge, this version and its extensions are the only off-the-shelf methods for capturing a notion of data distribution in high-dimensional spaces (density estimation techniques are usually not effective in high-dimensional spaces because of the curse of dimensionality (Scott, 2008)) and is applied to anomaly detection, among other tasks.

With a COOL network, one-class neural networks can be constructed in a similar manner by including one (or more) instances that are not within the one single class. In other words, following the same notation introduced in the previous section, the member (or members) of the second class are drawn from $D(x|y \notin \Theta)$ where $|\Theta| = 1$. Algorithm 3 summarizes the procedure.

Algorithm 3: The one-class neural network algorithm.

Input: The underlying probability distribution of the non-members of the single class $D'(X) = D(x|y \notin \Theta)$, number of samples from second class k, the set of available instances T
Output: COOL neural network N
step 1: form set α by sampling k instances from $D'(X)$
step 2: resize α to match the size of T by sampling uniformly from α
step 3: $T \leftarrow T \bigcup \alpha$
step 4: train N on T

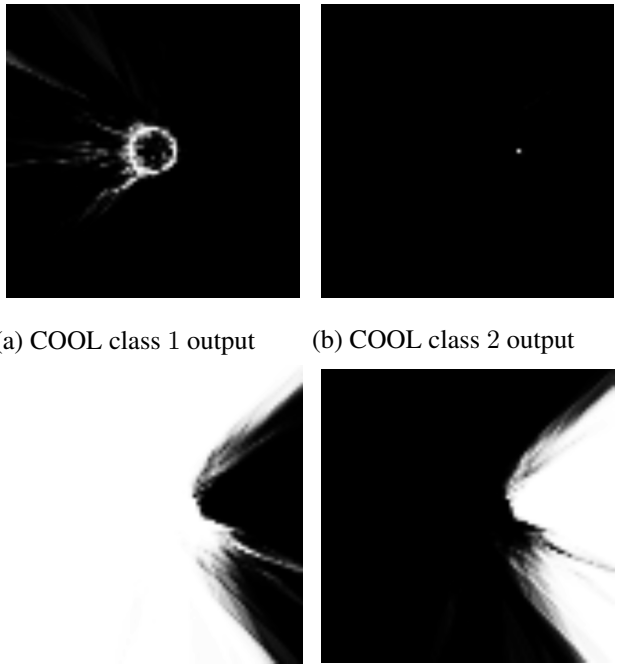

(a) COOL class 1 output (b) COOL class 2 output

(c) Regular neural network class 1 output (d) Regular neural network class 2 output

Figure 12: **One-class neural network visualizations.** Visualizations (a) and (b) depict the activation maps of the output aggregates on a one-class version of the visualization problem. These plots show how accurately one-class neural networks can capture the distribution of data in the only present class (a). For comparison, (c) and (d) depict the same activation maps when the COOL component in the one-class neural network is replaced with conventional neural network outputs.

Experiments

The experiments that follow demonstrate the abilities of one-class neural networks. In the first experiment a similar problem to the two-circle problem in visualization experiments section is introduced. More specifically, to form the training set, points are sampled uniformly within a circle C. To form an instance from the second class, an arbitrary point not within C is selected. Figure 12 depicts the activation map of the two output aggregates (DOO $= 5$) assigned to class one and two, respectively. For comparison, the result of the same experiment is also shown when typical neural networks are applied instead of COOL ones. These plots support the validity of the proposed one-class neural network in low-dimensional spaces where the regular neural network fails.

Next, MNIST tests the performance of one-class neural networks in a high-dimensional problem. In this experiment ten separate one-class neural networks are trained to learn the ten different classes of MNIST. Furthermore, one image with random pixel values is supplied as the only instance of the second class for all of these ten models. That way, each model is supposed to capture the data distribution of a single digit. Later, to test instance x, it is fed to all ten models and then assigned to the model with the maximum probability of inclusion (i.e. highest activation). In this manner, each model is treated as an unnormalized probability distribution function.

The main results are average test accuracies of 75.7% and 60% over five runs achieved by COOL (DOO $= 80$) and the conventional CNN, respectively. However, because classification accuracy is solely based on the maximum activation (probability of inclusion), it is not a good representative of the magnitude of probability values. Furthermore, in one-class learning problem the performance of the model is based on both its correct detection *and* rejection, which can be easily violated by classic discriminative model (due to the high data imbalance). Thus, to better demonstrate the difference between COOL and conventional CNNs in this experiment, figure 13 depicts the average accuracy and rejection rate of COOL versus regular CNNs over different epochs when 0.5 activation is set as the threshold for rejec-

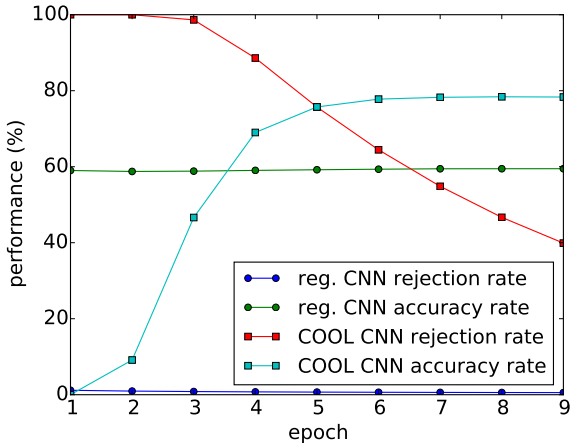

Figure 13: **Average accuracy vs. rejection rate in one-class experiment.** These plots show the average accuracy and rejection rates of the ten models over the test instances (from all 10 classes) in the first ten epochs of the MNIST one-class experiment when 0.5 is applied as the threshold value. The plots show that though regular CNNs reach some level of accuracy, their correct rejection rate is too low to make them an interesting one-class model.

tion (and minimum required activation in calculating accuracy). In other words, if a model's activation for test instance x is below 0.5, it is interpreted as rejecting the inclusion of x to the corresponding class (the class associated with the learning model). Notice in figure 13 that the rejection rate of the regular CNN is flat near the implausible value of zero, while that of COOL drops to about 60% when its accuracy rate flattens. Interestingly, Jain, Scheirer, and Boult (2014) report in their supplemental material an accuracy of about 52% for one-class RBF SVM in a similar experiment conducted on 6 classes of MNIST.[5] These results further affirm the advantage of COOL networks and their ability to capture the underlying distribution of data. Moreover, they support the effectiveness of the proposed one-class neural networks.

CIFAR-10 and CIFAR-100 Experiments

So far we have focused on the ability of COOL to capture the underlying distribution of data and generalize adequately. Another potential advantage of COOL is its ability to train larger architectures. As noted earlier, COOL in effect trains an exponential number of competing models together. An intriguing consequence is that it is possible to train very large architectures without overfitting. As results in this section will show, this property alone can sometimes lead to performance improvements.

Another important question is the rate of convergence in COOL. Is the aggressive competition in the output layer of these networks a hindrance to convergence, or, conversely, does it lead to more robust updates of parameters in the training phase and therefore accelerated learning?

In this section we conduct experiments with different network architectures on the CIFAR-10 and CIFAR-100 datasets to evaluate the effect of COOL on the rate of convergence and on performance. Using Torch[6], unlike previous experiments, the architectures in this section apply the latest techniques in deep learning to check their synergy with COOL, though again the learning procedure is kept as simple as possible.

CIFAR-10 Experiment

In this experiment a CNN with five convolutional layers followed by two fully-connected layers is trained on the first 45,000 training instances of the CIFAR-10 dataset. Each hidden layer is followed by a batch normalization layer and ReLU is applied as activation functions in all layers. More details are in the Appendix.

Figure 14 depicts the average classification accuracy over ten runs of the COOL network (with DOO $= 20$) versus ten runs of a conventional network with the same architecture, except for the last layer. The main result is that COOL networks converge more than two times faster while performing significantly better (t-test produces a p-value< 0.001). Preliminary experiments with a variety of different architectures yielded similar results.

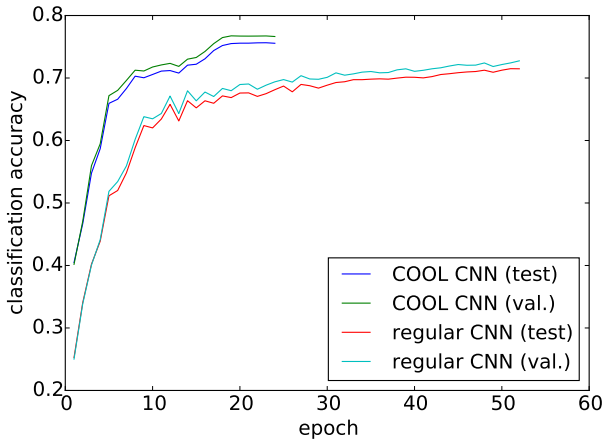

Figure 14: **Average classification accuracy of COOL and a conventional CNN with the same architecture on CIFAR-10 over** 10 **runs.** COOL speeds up learning significantly while also improving test accuracy. This result suggest more robust gradient updates during training.

CIFAR-100 Experiment

In this experiment a more comprehensive architecture is applied to the CIFAR-100 dataset. The unique characteristics of this experiment are the application of dropout and

[5] Jain, Scheirer, and Boult (2014) did not include the rejection rate in their report.

[6] `http://torch.ch/`

a light augmentation of the dataset according to the same procedure as in Clevert, Unterthiner, and Hochreiter (2016), where each image is padded by four zero pixels at all borders, randomly cropped to extract a 32×32 image, and also randomly horizontally flipped. To keep the training procedure simple, the dropout rate is kept constant and no momentum or weight decay is applied. More details are in the Appendix.

The classification accuracy of the conventional versus the COOL neural network without dropout over different epochs of the run are shown in figure 15. These results suggest, as expected, that dropout slows down the learning process significantly while improving performance. However, they also show once again that COOL converges faster and to a higher level of accuracy. For comparison, Table 1 lists some of the state-of-the-art results on CIFAR-100, including AlexNet (Krizhevsky, Sutskever, and Hinton, 2012), DSN (Lee et al., 2015), NiN (Lin, Chen, and Yan, 2014), Maxout (Goodfellow et al., 2013), All-CNN (Springenberg et al., 2014), Highway Network (Srivastava, Greff, and Schmidhuber, 2015), Fractional Max-Pooling (Graham, 2014), and ELU-Network (Clevert, Unterthiner, and Hochreiter, 2016).

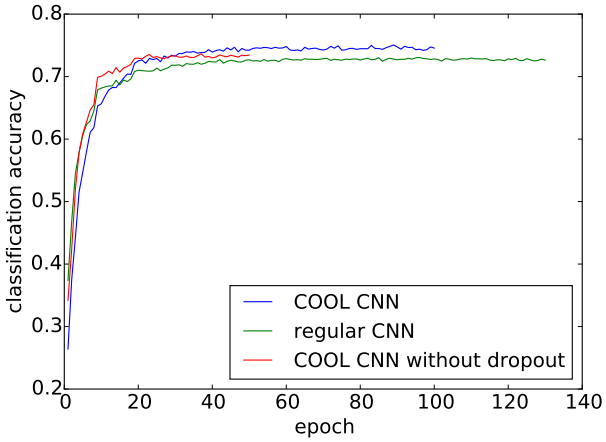

Figure 15: **Performance of COOL and non-COOL CNNs with the same architecture on CIFAR-100.** These plots show the synergy of COOL with dropout.

Discussion and Future Work

Formulating the notion of generalized classification helps to illuminate the rigidness of more conventional classification, which plays a central role in many of the most impactful benchmark results in deep learning. Generalized classification is arguably more closely aligned with a natural setting in which the set of classes or concepts is fluid and subject to expansion. It also enshrines the notion of the unknown as a fundamental ingredient in the framework of learning, an intuitive idea for humans where deep networks nevertheless can falter (as in fooling). Fitted learning thereby becomes a natural response to the problem of generalized classification, lighting a path towards new approaches and architectures.

Table 1: **Comparison of COOL network and some recent successful CNN architectures on CIFAR-100.**

Network	CIFAR-100 (test accuracy %)
AlexNet	54.2
DSN	65.43
NiN	64.32
Maxout	61.43
All-CNN	66.29
Highway Network	67.76
Fract. Max-Pooling	72.38
ELU-Network	**75.72**
COOL CNN w/o dropout	73.5
COOL CNN with dropout	74.72
CNN with dropout (with same arch. as COOL CNN)	72.79

One such new approach is the COOL network introduced in this paper. While the introductory experiments disclosed here are only an initial hint and more investigation will be necessary, by taking a significant step towards better fitted learning, COOL shows that the mechanisms that can enable a more dynamic form of learning are not necessarily complex or intricate. Rather, overrepresentation in the output layer combined with competition precipitates a partitioning of the input space among multiple overlapping models simultaneously, yielding a better fit to the data in aggregate. In this way, COOL in effect produces an ensemble of an exponential number of learning models, which is reminiscent of dropout. Even outside of separable learning and one-class neural networks, from the results in CIFAR-10 and CIFAR-100, the advantages of this modification seem to extend to conventional classification. It is an interesting question whether the main insights in COOL might bolster the case against single-neuron concepts in general, including in the brain.

At a practical level, the ability to add new classes to a model without the need to retrain opens the door to a new class of continual lifetime learners that expand their repertoire of concepts indefinitely. In this way, all the power of deep learning can be naturally brought to bear on lifetime learning, an exciting proposition. While a simple form of such continual learning is demonstrated in this paper where different modules learn different classes independently before later being combined, more sophisticated variations on this idea are conceivable and likely to emerge in the near future. For example, as new modules that address new concepts are added they can also potentially share hidden layers with old modules, allowing a safe form of expansion that also builds on accumulated knowledge. Tuning such hidden knowledge as the aggregate learner expands also offers a breadth of possible investigations. Furthermore, the potential benefits to recurrent neural networks remain unexplored yet promising.

In short, the fitted learning of COOL creates new possibilities and mitigates awkward problems (such as fooling) because it takes a qualitative step towards a more natural

kind of learning. That is, COOL is not just about more accurate classification, but about a different *perspective* on classification, towards generalized classification. While progress sometimes follows a path of quantitatively increasing performance, such as recent advances in accuracy in deep learning, often it also requires a new perspective to reimagine the status quo in a new light. While the results from COOL in this paper are too preliminary to declare victory in the problem of generalized classification, our hope is that COOL will provide such a productive shift in perspective that will at least set us down that path, and one day yield learners more in the dynamic and continual spirit of humans.

Acknowledgments

Special thanks to Omid Kardan for helpful discussions and comments.

References

Clevert, D.; Unterthiner, T.; and Hochreiter, S. 2016. Fast and accurate deep network learning by exponential linear units (ELUs). In *International Conference on Learning Representations*.

Duda, R. O.; Hart, P. E.; and Stork, D. G. 2012. *Pattern Classification*. John Wiley & Sons.

Fisher, R. A. 1936. The use of multiple measurements in taxonomic problems. *Annals of Eugenics* 7(2):179–188.

Gazzaniga, M. S. 2004. *The cognitive neurosciences*. MIT press.

Gershkoff-Stowe, L. 2002. Object naming, vocabulary growth, and the development of word retrieval abilities. *Journal of Memory and Language* 46(4):665–687.

Glorot, X., and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, volume 9, 249–256.

Goodfellow, I. J.; Warde-Farley, D.; Mirza, M.; Courville, A. C.; and Bengio, Y. 2013. Maxout networks. *ICML (3)* 28:1319–1327.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2672–2680.

Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and harnessing adversarial examples. In *13th International Conference on Document Analysis and Recognition, ICDAR 2015, Nancy, France, August 23-26, 2015*.

Graham, B. 2014. Fractional max-pooling. *arXiv preprint arXiv:1412.6071*.

Gross, C. G. 2002. Genealogy of the "grandmother cell". *The Neuroscientist* 8(5):512–518.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.

Jain, L. P.; Scheirer, W. J.; and Boult, T. E. 2014. Multi-class open set recognition using probability of inclusion. In *European Conference on Computer Vision*, 393–409. Springer.

Krizhevsky, A., and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical report, University of Toronto.

Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.

LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE* 86(11):2278–2324.

Lee, C.-Y.; Xie, S.; Gallagher, P.; Zhang, Z.; and Tu, Z. 2015. Deeply-supervised nets. In *AISTATS*.

Lin, M.; Chen, Q.; and Yan, S. 2014. Network in network. In *International Conference on Learning Representations*.

Nguyen, A.; Yosinski, J.; and Clune, J. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 427–436. IEEE.

Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* 115(3):211–252.

Scheirer, W. J.; Rocha, A. d. R.; Sapkota, A.; and Boult, T. E. 2013. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(7):1757 – 1772.

Schölkopf, B.; Platt, J. C.; Shawe-Taylor, J. C.; Smola, A. J.; and Williamson, R. C. 2001. Estimating the support of a high-dimensional distribution. *Neural Computation* 13(7):1443–1471.

Scott, D. W. 2008. The curse of dimensionality and dimension reduction. *Multivariate Density Estimation: Theory, Practice, and Visualization* 195–217.

Springenberg, J. T.; Dosovitskiy, A.; Brox, T.; and Riedmiller, M. 2014. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 15(1):1929–1958.

Srivastava, R. K.; Greff, K.; and Schmidhuber, J. 2015. Training very deep networks. In *Advances in neural information processing systems*, 2377–2385.

Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2013. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*.

Vapnik, V. 2013. *The nature of statistical learning theory*. Springer Science & Business Media.

Appendix

This appendix gives a more detailed description of the experiments in the paper. The COOL architecture introduces two new hyperparameters, namely the DOO and softness. The softness is always 1.0 except in the single visualization experiment that is designed to show its effect on the decision border. The DOO is typically 5, 10, or 20 depending on the type of dataset and problem because it is basically the number of input space partitions. For example, the two-circle problem with only two-dimensions does not in effect ask for as many partitions by COOL member units as in MNIST with 28×28 inputs. For tasks such as the one-class neural network a higher value makes sense because of the high data imbalance (and therefore tendency to overgeneralize) but in general we do not optimize this parameter systematically. Most experiments are aimed at simple architectures, except for CIFAR-10 and CIFAR-100, where state-of-the-art architectures are tested.

Visualization Experiments

The two-circle problem consists of 1,258 points uniformly sampled from two concentric circles with different radii. Table 2 lists the different feed-forward architectures applied throughout this paper to generate the visualizations.

Table 2: **The architectures applied in the visualization experiments.**

Experiment	Network Architecture
figure 3a	2,200,2
figure 3b	2,400,300,200,2
figure 4a-b	2,400,300,10
figure 4c-d	2.400,300,2
figure 5a	2,400,300,20
figure 5b	2,400,300,4
figure 6a	2,200,10
figure 6b	2,400,350,300,250,200,10
figure 8	2,800,300,100,15
figure 9	2,600,200,10
figure 12a-b	2,600,200,10
figure 12c-d	2,600,200,2

Experiments with MNIST

The MNIST dataset is a popular machine learning benchmark consisting of 60,000 training and 10,000 28×28 digit test images. In all the experiments the first 50,000 training images are used as the training set and the rest as validation. The neural network architectures are kept simple (most CNN architectures are inspired by LeNet-5 (LeCun et al., 1998)) and all the hidden units apply a saturating activation function (logistic). In the preprocessing step we simply transform all pixel intensity values to the $[0, 1]$ interval by dividing them by 255. Finally, the learning rate is constant and the weight initialization follows Glorot and Bengio (2010).

Fooling Experiment In this experiment we first train two CNNs with a $[Conv(1 \rightarrow 20, 5 \times 5), maxPooling(2 \times 2), Conv(20 \rightarrow 50, 5 \times 5), maxPooling(2 \times 2), fullyConnected(50 \rightarrow 400), fullyConnected(400 \rightarrow 10 \times DOO)]$ architecture, where the DOO is 1 for the regular and 10 for the COOL network, respectively. The learning rate is 0.001 and the training continues for at most 200 epochs.

Next are the trials on both trained networks. Each trial consists of (1) generating a 28×28 image with pixels sampled from the uniform distribution in the $[0, 1]$ interval, (2) training a FGN (learning rate $= 0.00001$) consisting of a single fully-connected layer that transforms this image into a 28×28 fooling image with all pixels in the $[0, 1]$ interval. In the case of the COOL network, the FGN is trained based on the errors backpropagating from member units of a specific class and not the final result after multiplication. As a side note we also tried FGNs with several layers but the single-layer FGN works the fastest while serving the role of constraining the pixel intensities to the valid range of $[0, 1]$.

Separable Concept Learning Experiment This experiment evaluates the performance of COOL versus a regular neural network on both CNN and MLP architectures. In the training phase, five models are trained, each on different pairs of classes. During testing, a test instance is fed to all the five models and the class with the highest activation is selected as the final decision. The validation set includes examples from all the 10 classes and is the only means to tune the amount of generalization in the final models.

The MLP architecture consists of 3 hidden layers each with 500 units and the DOO is one and 10 for the regular and COOL network, respectively. The training applies a learning rate of 0.005 and is continued for the maximum of 120 and 250 epochs for COOL and regular network, respectively, and the reported MLP results are the average of 2 runs.

The CNN architecture is $[Conv(1 \rightarrow 20, 5 \times 5), maxPooling(2 \times 2), Conv(20 \rightarrow 50(10), 5 \times 5), maxPooling(2 \times 2), fullyConnected(50 \rightarrow 400), fullyConnected(400 \rightarrow 10 \times DOO)]$, where to speed up training we apply strides of 2 to both convolutional layers and the second convolutional layer is mapped randomly to half of the feature maps from the previous layer. Again, the DOO is set to 1 and 10 for conventional and COOL network, respectively, and the learning rate is 0.005. The CNN results are the average of 10 runs of this experiment.

One-class Neural Network Experiment In these experiments the same CNN architecture as in the separable concept learning experiment is applied to train 10 different models. The DOO for the regular and COOL networks are one and 80 respectively. We also tried DOOs of 20 and 50 but found that higher DOO leads to better results in the trade-off between classification accuracy and rejection rate.

Experiments on CIFAR-10

The CIFAR-10 dataset consists of 50,000 of training and 10,000 testing $3 \times 32 \times 32$ color images from 10 categories of some animal and transportation media. The validation

set is the last $5{,}000$ images of the training images, which are excluded from the training set. For preprocessing, we simply subtract the mean and divide the standard deviation of each pixel location from pixel values. The architecture is: $[Conv(3 \rightarrow 16, 3 \times 3), Conv(16 \rightarrow 32, 5 \times 5), maxPooling(2 \times 2), Conv(32 \rightarrow 64, 3 \times 3), Conv(64 \rightarrow 128, 5 \times 5), maxPooling(2 \times 2), Conv(128 \rightarrow 256, 3 \times 3), fullyConnected(256 \rightarrow 800), fullyConnected(800 \rightarrow 10 \times DOO)]$,
where all layers (except for the output and pooling layers) follow batch normalization and apply ReLU as the activation function. The DOO is 1 and 10 for regular and COOL networks, respectively and the learning rate is 0.01 with the batch sizes of 500.

Experiments on CIFAR-100

The CIFAR-100 dataset is another dataset of small images similar to CIFAR-10 but it contains 100 categories with the same total number of images. The experiment on CIFAR-100 applies the same preprocessing as with CIFAR-10 but on image derivatives (with respect to x) instead of the raw pixel values. The architecture is:
$[Conv(3 \rightarrow 128, 1 \times 1), Conv(128 \rightarrow 512, 3 \times 3), maxPooling(2 \times 2), Conv(512 \rightarrow 512, 1 \times 1), Conv(512 \rightarrow 1024, 3 \times 3), Dropout(0.1), maxPooling(2 \times 2), Conv(1024 \rightarrow 1024, 1 \times 1), Conv(1024 \rightarrow 3072, 3 \times 3), Dropout(0.2), maxPooling(2 \times 2), Conv(3072 \rightarrow 3072, 1 \times 1), Conv(3072 \rightarrow 8192, 3 \times 3), Dropout(0.3), maxPooling(2 \times 2), Conv(8192 \rightarrow 8192, 1 \times 1), Conv(8192 \rightarrow 100 \times DOO, 3 \times 3), maxPooling(2 \times 2)]$,
where each dropout layer and the last convolutional layer follow batch normalization and ReLU is the activation function. The DOO is 1 and 20 for regular and COOL networks, respectively and the learning rate is 0.001, which is halved every 10 epochs.